

AT9-99-274

PATENT

AN APPARATUS AND METHOD FOR DISTRIBUTING AND  
COLLECTING BULK DATA BETWEEN A LARGE NUMBER OF MACHINES

CROSS-REFERENCE TO RELATED APPLICATIONS

Sub  
A1

Related subject matter may be found in the following commonly assigned, co-pending U.S. Patent Applications, which are hereby incorporated by reference herein:

Serial No. \_\_\_\_\_ (AT9-99-275), entitled "APPARATUS FOR DATA DEPOTING AND METHOD THEREFOR";

Serial No. \_\_\_\_\_ (AT9-99-276), entitled "APPARATUS FOR RELIABLY RESTARTING INTERRUPTED DATA TRANSFER AT LAST SUCCESSFUL TRANSFER POINT AND METHOD THEREFOR";

Serial No. \_\_\_\_\_ (AT9-99-655), entitled "APPARATUS FOR CONNECTION MANAGEMENT AND METHOD THEREFOR" and filed concurrently herewith;

Serial No. \_\_\_\_\_ (AT9-99-324), entitled "COMPUTER NETWORK CONTROL SYSTEMS AND METHODS" and filed concurrently herewith;

Serial No. \_\_\_\_\_ (AT9-99-325), entitled "METHODS OF DISTRIBUTING DATA IN A COMPUTER NETWORK AND SYSTEMS USING THE SAME";

Sub  
A1  
Serial No. \_\_\_\_\_ (AT9-99-315), entitled "SYSTEMS AND METHODS  
FOR REAL TIME PROGRESS MONITORING IN A COMPUTER NETWORK;

Serial No. \_\_\_\_\_ (AT9-99-316), entitled "APPARATUS FOR  
AUTOMATICALLY GENERATING RESTORE PROCESS DURING SOFTWARE  
DEPLOYMENT AND METHOD THEREFOR"; and

Serial No. \_\_\_\_\_ (AT9-99-323), entitled "AN APPARATUS FOR  
JOURNALING DURING SOFTWARE DEPLOYMENT AND METHOD  
THEREFOR".

#### TECHNICAL FIELD

The present invention relates generally to data processing systems, and in particular, to bulk data distributions within networked data processing systems.

5

#### BACKGROUND INFORMATION

Present day data processing systems are often configured in large multi-user networks. Management of such networks may typically include the need to transfer large amounts data to an endpoint system from a source system (or, simply, "a source") and the collection of information, for example, error reports from a multiplicity of endpoints systems (or, simply, "endpoints").

10

Such large data transfers may occur within a network, for example, to distribute software updates. The system administrator may need to allocate a specific period of time for the bulk data transfer to more efficiently utilize network resources. This may typically occur when the communication load on the system is lowest, usually at night when most endpoint users are not working at their stations. The system administrator may load the bulk data and the corresponding transfer instructions onto the network system's source, or server, in preparation for the transfer. At the predetermined time set by the administrator, the server will push the data while ensuring that the bulk data is successfully transferred to each of the desired endpoint locations. However, during the transfer a portion of the system server is dedicated to the data transfer and thus unavailable for other networking tasks. Moreover, as the number of endpoints which must be simultaneously serviced by the bulk data distribution increases, network bandwidth demands are concomitantly increased. This complicates scalability of the bulk distribution systems.

Therefore, a need exists in the art for a bulk distribution mechanism that can transfer large amounts of data between network connected subsystems (or nodes) while maintaining scalability. Additionally, there is a need in such distribution mechanisms for methods and apparatus to distribute bulk data to a multiplicity of endpoints and to collect bulk data, including large log files, from the endpoints.

## SUMMARY OF THE INVENTION

The aforementioned needs are addressed by the present invention.

Accordingly, the present invention provides a general service that allows applications to asynchronously distribute large amounts of data from a source node, to multiple  
5 destination nodes, to process the data on each destination node, and then to collect the results on one or more "report-to" nodes.

The present invention includes fan-out nodes (which will also be referred to as repeaters) and methods therefor, which are nodes on the network which receive bulk data streams, and retransmit the data to the follow-on fan-out nodes or to endpoints.  
10 Additionally, the fan-out nodes receive bulk results from downstream and retransmit them to upstream fan-out nodes or final report-to nodes.

Additionally, the present invention includes a method and apparatus for enqueueing the distribution information received from a requesting application in a persistent queue at the repeaters according to its priority, and returns to the  
15 application and unique ID that can be used as a correlator for the results.

The invention also provides for a different maximum number of available sessions according to a predetermined set of selectable transmission priority levels. A distribution with a given priority level can use the number of sessions reserved for its priority level plus any sessions allocated for lower priority levels.

20 The foregoing has outlined rather broadly the features and technical advantages of the present invention in order that the detailed description of the

invention that follows may be better understood. Additional features and advantages of the invention will be described hereinafter which form the subject of the claims of the invention. It is important to note the drawings are not intended to represent the only form of the invention.

## BRIEF DESCRIPTION OF THE DRAWINGS

For a more complete understanding of the present invention, and the advantages thereof, reference is now made to the following descriptions taken in conjunction with the accompanying drawings, in which:

5           FIGURE 1 illustrates, in block diagram form, a data processing network in accordance with one embodiment of the present invention;

          FIGURE 2 illustrates, in block diagram form, a data processing system implemented in accordance with an embodiment of the present invention;

10           FIGURE 3A illustrates, in flowchart form, a distribution request methodology in accordance with an embodiment of the present invention;

          FIGURE 3B illustrates, in tabular form, a distribution structure in accordance with an embodiment of the present invention;

          FIGURE 4A illustrates, in flowchart form, a methodology to transfer data in accordance with an embodiment of the present invention;

15           FIGURE 4B is a continuation of FIGURE 4A and illustrates, in flowchart form, a methodology to transfer data over a network in accordance with an embodiment of the present invention;

20           FIGURE 5 illustrates, in flowchart form, a methodology implemented to determine priority resource availability in accordance with an embodiment of the present invention; and

FIGURE 6 illustrates in flowchart form, a database management methodology implemented in accordance with an embodiment of the present invention.

## DETAILED DESCRIPTION

5 The present invention is a method and apparatus for distributing and collecting data between an originating source system and a plurality of endpoint systems (which may also be referred to as "endpoint nodes" or simply "endpoints"). The method and apparatus provide a general service that allows applications to asynchronously distribute large amounts of data from a source node to multiple destination nodes, to process the data on each destination node, and then to collect the results of that processing on one or more report-to nodes.

10 According to the principles of the present invention, the present invention has an originating source system followed by repeaters. The use of repeaters allows data to be delivered essentially simultaneously to a large number of machines. The present invention can be scaled to handle more destinations by adding repeaters. In the following description, numerous specific details are set forth to provide a thorough understanding of the present invention. However, it will be obvious to those skilled in the art that the present invention may be practiced without such specific details. In 15 other instances, well-known circuits have been shown in block diagram form in order not to obscure the present invention in unnecessary detail. For the most part, details concerning timing considerations and the like have been omitted inasmuch as such details are not necessary to obtain a complete understanding of the present invention and are within the skills of persons of ordinary skill in the relevant art. 20



A more detailed description of the implementation of the present invention will subsequently be provided. Prior to that discussion, an environment in which the present invention may be implemented will be described in greater detail.

FIGURE 1 illustrates a communications network 100. The subsequent discussion and description of FIGURE 1 are provided to illustrate an exemplary environment used by the present invention.

The network system 100 includes source system 101, one or more fan-out or, repeaters 110, 111, 118, 119, and a plurality of endpoints 112-117. Additionally, certain repeaters, such as 118 and 119, are directly connected to one or more endpoints, in the exemplary embodiment of FIGURE 1, endpoints 112-114 or 115-117, respectively, and may be referred to as "gateway" repeaters (or, simply, "gateways").

Source system 101 provides distribution services with respect to resources 112-117. Note that source system 101 and endpoints 112-117 interfaces to repeaters 110 and 111 using the same methodologies as repeaters 110 and 111 interface with, for example, repeaters 118 and 119. Viewed logically, source system 110 and endpoints 112-117 each may include a "repeater". In other words, as an artisan of ordinary skill would recognize, as used herein, a repeater may be a logical element, that may be, but is not necessarily associated with a physical stand-alone hardware device in network 100. Repeater 110 may be the primary repeater through which resources 112-114 receive their data transfers, and repeater 111, likewise, may

5 primarily service endpoints 115-117. Additionally, any report-back of successful transfers will be transmitted primarily via the endpoints primary domain except as explained below. It would be understood by an artisan of ordinary skill that additional repeaters may be inserted into the network and may be arranged in a multi-level hierarchy according to the demands imposed by the network size.

Gateway repeaters 118 and 119 are such repeaters in the exemplary embodiment of FIGURE 1.

10 However, network system 100 may provide cross connections in order to provide redundant, parallel communication paths should the primary communication path to the endpoint become unavailable. For example, in FIGURE 1, endpoint 114 has a primary pathway to source system 101 through repeaters 118 and 110. (A source system, such as source system 101 may also be referred to as a source node.) Should repeater 110 become unavailable, source system 101 can transfer bulk data to endpoint 114 via an alternative pathway through repeaters 118 and 111. Additionally, 15 should repeater 118 become unavailable, endpoint 114 may receive data via repeaters 111 and 119. Source system 101 maintains database 120 for storing information used in managing a data distribution. A methodology which may be used to process the information to be stored in database 120 will be described in conjunction with FIGURE 6.

20 Sub A3 Referring next to FIGURE 2, an example is shown of a data processing system 200 which may be used to implement a source system such as system 101,

5 repeaters, such as repeaters 110, 111, 118, or 119 or endpoints, such as endpoints 112-117, executing the methodology of the present invention. The system has a central processing unit (CPU) 210, which is coupled to various other components by system bus 212. Read only memory ("ROM") 216 is coupled to the system bus 212 and includes a basic input/output system ("BIOS") that controls certain basic functions of the data processing system 200. Random access memory ("RAM") 214, I/O adapter 218, and communications adapter 234 are also coupled to the system bus 212. I/O adapter 218 may be a small computer system interface ("SCSI") adapter that communicates with a disk storage device 220. Disk storage device 220 may be used to hold database 120, FIGURE 1. Communications adapter 234 interconnects bus 212 with the network as well as outside networks enabling the data processing system to communicate with other such systems. Input/Output devices are also connected to system bus 212 via user interface adapter 222 and display adapter 236. Keyboard 224, track ball 232, mouse 226 and speaker 228 are all interconnected to bus 212 via user interface adapter 222. Display monitor 238 is connected to system bus 212 by display adapter 236. In this manner, a user is capable of inputting to the system throughout the keyboard 224, trackball 232 or mouse 226 and receiving output from the system via speaker 228 and display 238.

20 Preferred implementations of the invention include implementations as a computer system programmed to execute the method or methods described herein, and as a computer program product. According to the computer system

implementation, sets of instructions for executing the method or methods are resident in the random access memory 214 of one or more computer systems configured generally as described above. Until required by the computer system, the set of instructions may be stored as a computer program product in another computer memory, for example, in disk drive 220 (which may include a removable memory such as an optical disk or floppy disk for eventual use in the disk drive 220). Further, the computer program product can also be stored at another computer and transmitted when desired to the user's work station by a network or by an external network such as the Internet. One skilled in the art would appreciate that the physical storage of the sets of instructions physically changes the medium upon which it is stored so that the medium carries computer readable information. The change may be electrical, magnetic, chemical, biological, or some other physical change. While it is convenient to describe the invention in terms of instructions, symbols, characters, or the like, the reader should remember that all of these and similar terms should be associated with the appropriate physical elements.

Note that the invention may describe terms such as comparing, validating, selecting, identifying, or other terms that could be associated with a human operator. However, for at least a number of the operations described herein which form part of at least one of the embodiments, no action by a human operator is desirable. The operations described are, in large part, machine operations processing electrical signals to generate other electrical signals.

Refer now to FIGURE 3A in which is illustrated a flow chart of a methodology 300 for processing distribution requests in accordance with an embodiment of the present invention.

Applications resident on endpoint systems may request a distribution of bulk data. The application may, for example, maintain a periodic update schedule, or respond to a user update request. When an application requires updating, or the application otherwise requires a distribution of data, the application requests the data from the source system, step 301. The source system may be a server, such as server 110, FIGURE 1.

Applications may submit a request by involving a predetermined method which is included in each repeater. In step 302, the request may be transmitted to the source system via one or more repeaters in similar fashion. The distribution request may be encapsulated in a data structure, which may be passed to the method. A data structure 315 which may be used in the present invention is schematically shown in FIGURE 3B. Data structure 315 includes entries 317, 319, 321 and 323. These entries respectively include a distribution data specifier which identifies the data to be distributed in response to the request and the location of the data, a list of destination node identifiers (IDs) that specify the endpoints that are to receive the data, the method on the endpoint that will receive and process the data, and the method that will receive and process results information from each endpoint node receiving the distribution. The method identifier in field 321 informs the source system of the

method on the endpoint system to be invoked to receive and process the data. As described further below, endpoints provide result information to the source system. The method identifier in field 323 informs the source system of the method on the source system that a repeater will invoke to receive and process the result information sent by the endpoint. In an embodiment of the present invention, this method may be implemented in accordance with the Common Object Request Broker Architecture (CORBA). The CORBA is specified in "The Common Object Request Broker: Architecture and Specification," Version 2.3, June 1999, which is hereby incorporated herein by reference.

In step 320, the request is received by the source system, and, in step 340, the source system enqueues the distribution information from the distribution structure in a database, such as database 120, FIGURE 1. The distribution information is enqueued in accordance with a preselected distribution priority which may, in an embodiment of the present invention, be one of three levels: high, medium, or low. The use of the priority schedule in transferring data is discussed in detail in conjunction with FIGURES 4A, 4B and 5.

The source system provides the target endpoint(s) a distribution identification (DID) in step 350. The DID is used by the endpoints receiving the distribution to tag the results information, whereby the source system may correlate the results information when it is received from the endpoints. Additionally, the DID is used by the endpoints to check the data transmission as it is occurring as discussed in more

detail in the co-pending commonly-owned U.S. Patent Application entitled "Apparatus for Restarting Interrupted Data Transfer and Method Therefor," incorporated herein by reference. The application can exit to allow the server and endpoints to perform other operations, step 370. The supplied notification method will be called by the source system to receive results information sent from each endpoint node. The distribution of the data is described in FIGURE 4A and FIGURE 4B.

Refer now to FIGURE 4A and FIGURE 4B in which is illustrated a flow chart of methodology 400 for distributing and collecting bulk data between data processing systems. Methodology 400 may be implemented by a data processing network such as network 100, FIGURE 1.

In step 401, data for distribution is loaded on to the source system, for example source system 101, FIGURE 1. The administrator may allocate the appropriate endpoints to receive the data. Or, as described above, the appropriate endpoints may request the distribution based on applications running on the endpoints. Status information regarding the distribution is retained in a database which, in an embodiment of the present invention, may be maintained by the source system, such as source system 101, FIGURE 1, as previously described. In step 402, a database entry corresponding to the distribution loaded in step 401 is created. If the system administrator has identified target endpoints, these are stored in the database entry, which will incorporate status information for each distribution endpoint.

Additionally, endpoint destinations identified in distribution requests as described above in conjunction with FIGURE 3 will also be incorporated in the database entry for the distribution. In an embodiment of the present invention, employing a network such as network 100, FIGURE 1, the database entry may be included in database 120 maintained by source system 101.

In step 404, a connection is opened to a target repeater. The target repeater may be a gateway repeater, as discussed below in conjunction with step 410, or may be an intermediate repeater for fanning out the distribution to one or more gateways. In step 405, it is determined if a session is available. Network bandwidth management is implemented by allocating resources (referred to as sessions) for transferring data in accordance with a priority scheme in which a particular distribution is assigned one of a predetermined set of priority levels. A method and apparatus for connection management is described in detail in the corresponding commonly owned U.S. Patent Application entitled "An Apparatus for Connection Management and Method Therefor," incorporated herein by reference. For each priority level a predetermined "pool" of resources, or sessions, for transferring data, is allocated and a distribution may use a session from the pool corresponding to the priority level of the distribution or from a lower priority level pool. That is, high priority pool may have  $j$  sessions allocated, a medium priority pool may have  $b$  sessions and a low priority pool may have  $l$  sessions. The availability of a session is determined based on the available bandwidth and the priority level of the distribution



as described in detail in conjunction with FIGURE 5. If in step 405 a connection is available, while the connection is established, step 411, data is transferred to the repeater, in step 406. The transfer data is stored on the repeater, step 408. The storage may be temporary or permanent in accordance with an explicit administrator command, or control information provided by an application requesting the distribution. The storage, or depoting, of data on a repeater is described in detail in the co-pending, commonly owned U.S. Patent Application entitled, "Apparatus for Data Depoting and Method Therefor," ~~(Attorney Docket No. AT9-99-275)~~, and incorporated herein by reference.

In step 415, it is determined if all of the data constituting the distribution has been transferred. If not, methodology 400 returns to step 411, and data transfer continues by following the "True" branch to step 406, provided the connection has not failed.

If the connection has failed, methodology 400 returns to step 405. Methodology 400 then loops between steps 405 and 407, where, in step 407 it is determined if a preselected amount of time has elapsed until, in step 405, the connection has become available.

If, in step 407, the preselected time interval has elapsed, which may be referred to as "timing out," it is then determined, in step 425, if a preselected distribution lifetime has expired. If, in step 425, the distribution lifetime has not expired, then methodology 400 proceeds to step 404 to open a connection to an

alternative repeater, wherein methodology 400 then performs step 405 to determine the availability of a connection to the alternative repeater, in the same fashion as discussed above. Recall, as discussed hereinabove in conjunction with FIGURE 1, a network in accordance with an embodiment of the present invention may include a plurality of parallel repeater paths between the distribution source system and the target endpoints. However, in an alternative embodiment of the present invention, the network, may not implement parallel paths, and in such an alternative embodiment, in steps 404 and 405, the connection is retried to the same target.

A method and apparatus for connection management which may be used to implement steps 404-415 is described in the commonly owned, co-pending U.S. Patent Application entitled "Apparatus for Connection Management and Method Therefor" (~~Attorney Docket No. AT9-99-655~~) incorporated herein by reference.

If, however, in step 425 the distribution lifetime has expired, the distribution is aborted, step 419. In step 420, the status of the distribution, which is included in the results information associated with the distribution, is returned to the source system, or other preselected "report-to" systems which may include an endpoint system running the application requesting the distribution.

Returning to step 405, if a connection is reestablished before the timeout period has elapsed, methodology 400 proceeds by the "Yes" branch to step 411 and then to step 406 to transfer additional data. The data transfer may resume at a preselected checkpoint wherein the storage of data, in step 408, is periodically

committed to a nonvolatile storage medium, for example a disk storage device 220,  
FIGURE 2. The transfer of data, in step 406, may then resume subsequent to the last  
permanently stored data. Such a data transfer using preselected checkpoints to  
resume interrupted data transfers is described in detail in the commonly owned,  
co-pending U.S. Patent Application entitled, "Apparatus for Reliably Restarting  
Interrupted Data Transfer and Method Therefor" (~~Attorney Docket No. AT9-99-276~~),  
incorporated herein by reference.

Methodology 400 loops between steps 404-408, 411, 415, and 425 until the  
full distribution is transferred, or the distribution aborts in step 420. When the  
transfer completes as determined in step 415, methodology 400 then proceeds to  
step 409 and returns status information to one or more "report-to" machines as  
previously discussed in conjunction with step 420. The status may be sent to the  
endpoint requesting a distribution using the DID provided as described in conjunction  
with step 315, FIGURE 3, above.

In step 410 it is determined if the current repeater is a gateway. If so, in  
step 412 a connection to the endpoints receiving the distribution is established. In  
step 414, it is determined if the connection is available. Again, as discussed  
hereinabove in conjunction with step 405, a preselected number of connections may  
be available in accordance with a priority scheme. If a connection is not available,  
methodology 400 proceeds through steps 414 and 418 until a connection is  
established, or in step 418, the distribution lifetime expires, wherein the distribution

aborts in step 419, and the distribution status is returned in step 420, as previously described.

5 If, however, in step 414, a connection at the requested priority is available, while the connection is established, in step 417, data is transferred to the endpoint, step 416. Transfer continues, while, in step 422, the complete distribution has not been transferred, in which case the methodology 400 loops between steps 414, 417, 418, 416 and 422. The aforementioned method and apparatus for connection management described is the co-pending, commonly owned U.S. Patent Application entitled "Apparatus for Connection Management and Method Therefor, incorporated herein by reference, may be used to implement steps 412-418.

10 The data transfer in steps 416 may occur using the checkpoint process previously discussed in conjunction with the data transfer between repeaters, step 406, and described in detail in the co-pending, commonly owned U.S. Patent Applicant entitled "Apparatus for Data Depoting and Method Therefor," ~~(Attorney~~  
15 ~~Docket No. AT9-99-275)~~ incorporated herein by reference. On completion of the transfer, step 422, status is returned in step 420 to one or more report-to systems.

20 Methodology 400 may also be used to return distribution results information to the report-to systems. In using methodology 400 in this way, the report-to systems play the role of endpoint systems with respect to data distributions, and each endpoint receiving or repeater relaying the distribution data play the role of source systems. In this way, for example, log files generated by an installation program may be returned

to a preselected report-to system. Otherwise, as would be recognized by an artisan of ordinary skill, FIGURE 4 is unchanged.

A methodology for opening connections in accordance with a preselected distribution priority level which may be used in conjunction with steps 404 and 405, and 412 and 414 of methodology 400 will now be described. FIGURE 5 illustrates, in flowchart form, a methodology 500 for performing a session availability determination in accordance with an embodiment of the present invention.

In step 505, a session request is received, for example, from methodology 400, FIGURE 4 when opening a connection. A particular distribution in an embodiment of the present invention may be assigned one of three priority levels, low, medium or high, which determines the order in which the distribution is handled by a repeater. Distributions with higher priority levels are handled before those with lower priority, and distributions with the same priority level are handled in the order in which they are received by the repeater. The priority level may be set by an application requesting the distribution. A default priority may be set at the source repeater when it receives a distribution request from an application.

In step 518, it is determined if the distribution has a high priority level. If not, then in step 520, it is determined if the distribution has a medium priority level. If not, then the distribution has a low priority, step 530 and, in step 535, it is determined if a session is available in the low-priority pool. If low priority session is available, then in step 550, methodology 500 signals that a connection is available. In an

embodiment of the present invention in accordance with methodology 400, FIGURES 4A and 4B, the information from step 550 may be received in steps 405 and 414 in response to the opening of connections in steps 404 and 412, respectively. Conversely, if no low priority sessions are available in step 535, in step 540 methodology 500 signals that no session is available.

Returning to step 510, if the data transfer is determined to be a high priority transfer, then in step 515 it is determined if a high priority session is available. If so, then methodology 500 proceeds to step 550. Otherwise, if high priority sessions are unavailable, that is, fully used by other distributions, then in step 525 it is determined if a medium priority session is available. Again, if a medium priority session is available, then step 550 is performed; otherwise, in step 535, it is determined if a low priority level session is available. If so, step 535, then step 550 is performed; otherwise, no sessions are available and methodology 500 proceeds to step 540.

Similarly, if in step 510, it has been determined that the distribution is not a high priority distribution, it is determined if the distribution has a medium priority, step 520. If not, it must again be a low priority distribution, step 530, previously described. Otherwise, in step 520 it is a medium priority distribution, and in step 525 it is determined if a medium priority session is available. As before, if no medium priority sessions are available, it is determined if a low priority session is available, step 535. In this manner, a data distribution with a given priority level can use the number of sessions reserved for its priority level plus any sessions allocated to lower

priority levels. If no sessions are available at the assigned priority or lower, the methodology signals no available sessions, step 540, as previously described.

Methodology 500 overcomes the tasking conflicts that might arise when a repeater is processing a lower priority distribution and a higher priority distribution is received. Without the ability to specify priority sessions, the high priority distribution would have to wait until the lower priority distribution was complete, or alternatively, the lower priority distribution be interrupted, causing inefficiencies within the distribution system, by having to subsequently re-distribute the lower priority distribution.

As discussed hereinabove, in conjunction with, for example, FIGURE 4, results reports are generated by repeaters and endpoints reporting status of the transfer of a data distribution. Results are sent to one or more "report-to" machines, which may include the distribution source system. Results are stored in the distribution database, such as database 120 in FIGURE 1, maintained by the distribution manager, and may be entered into the database in conjunction with methodology 600, FIGURE 6. In step 601, methodology 600 loops until results are received. Results received are correlated in step 610 using the DID, which may be provided in accordance with methodology 300 for processing distribution requests, FIGURE 3. Recall that distributions may be sent to a plurality of endpoints, each of which may be reached via a different sequence of repeaters, as discussed hereinabove in conjunction with FIGURE 1. Thus, results for a particular distribution may be forwarded to the

"report-to" machines from several endpoints at different times. The DID, which is unique to the particular distribution, then allows the results associated with that distribution to be correlated for updating the database entry corresponding to the distribution, in step 630. In this way, the distribution manager maintains a current status for each distribution.

To avoid indefinitely filling the database, distribution records can be deleted. Automatic deletion may be provided after the expiration of a preselected record lifetime. In step 635, if the entry lifetime has expired, the corresponding entry is deleted in step 650. Even if, however, the entry lifetime has not expired, the entry may be deleted manually. In step 640, it is determined if a system administrator has initiated a command to delete an entry. If so, the entry is, again, deleted in step 650. Otherwise, methodology 600 returns to step 601 to continue to receive distribution results.

The foregoing has outlined rather broadly the features and technical advantages of the present invention in order that the detailed description of the invention that follows may be better understood. Additional features and advantages of the invention will be described hereinafter which form the subject of the claims of the invention.